



Welcome United States Patent and Trademark Office

Search Session History

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Wed, 1 Feb 2006, 12:33:38 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Search Query Display

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- #1(( data identification<in>metadata ) ) <AND> (client and server)
- #2graph traversal technique
- #3graph traversal technique
- #4(network cache<IN>metadata)
- #5(data defragmentation<IN>metadata)
- #6(data threshold<IN>metadata)
- #7(data inclusion<IN>metadata)
- #8(incomplete object<IN>metadata)
- #9(synchronous or synchronize or simultaneous<IN>metadata)
- #10((( data identification<in>metadata ) ) <AND> (client and server)) <AND> (graph traversal technique)
- #11(((synchronous or synchronize or simultaneous<IN>metadata)) <AND> ((( data identification<in>metadata ) ) <AND> (client and server)) <AND> (graph traversal technique))
- #12(((data defragmentation<IN>metadata)) <AND> (((synchronous or synchronize or simultaneous<IN>metadata)) <AND> ((( data identification<in>metadata ) ) <AND> (client and server)) <AND> (graph traversal technique)))



## EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	48	(James near4 Challenger).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L2	26	(Louis near4 Ralph).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L3	15	(Louis near2 Degenaro).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L4	16	(Robert near2 Filepp).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L5	40	(Arun near2 Kwangil near2 Iyengar).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L6	13	(Richard near2 Pervin near2 King).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L7	48	(James near4 Challenger).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L8	15	(Louis near2 Degenaro).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L9	16	(Robert near2 Filepp).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11

## EAST Search History

L10	40	(Arun near2 Kwangil near2 Iyengar).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L11	13	(Richard near2 Pervin near2 King).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L12	1	L7 and L8 and L9 and L10 and L11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L13	0	711//\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L14	27744	"711"//\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L15	6085	"709"/203.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L16	2	(incomplete adj version) near3 object\$2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L17	131596	client and server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L18	30460	data near2 construct\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L19	276	fragment near2 identifier	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11

## EAST Search History

L20	5237	fragment near2. (identifier or ID)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L21	82539	(data or object) adj2 position	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L22	343	(incomplete near2 object\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L23	131596	client and server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L24	30460	data near2 construct\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L25	3937	L23 and L24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L26	343	(incomplete near2 object\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L27	3937	L23 and L24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L28	0	L27 and L26	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L29	27744	"711"/\$.cls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11

## EAST Search History

L30	6085	"709"/203.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L31	33577	L29 or L30	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:11
L32	33577	L29 or L30	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:12
L33	32	L32 and L26	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:12
L34	32	L32 and L26	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:12
L35	3	L34 and L24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:12
L36	5237	fragment near2 (identifier or ID)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:12
L37	1	L36 and L34	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/02/01 12:12



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

+client, +server, +data +constructing, +fragment +identifier {

**SEARCH**



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

**[client](#) [server](#) [data](#) [constructing](#) [fragment](#) [identifier](#) [data](#) [object](#) [incomplete](#) [object](#)**

Found **458** of **169,866**

Sort results by



[Save results to a Binder](#)

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Display results



[Search Tips](#)

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 [Client-server computing in mobile environments](#)



Jin Jing, Abdelsalam Sumi Helal, Ahmed Elmagarmid

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

**Publisher:** ACM Press

Full text available: [pdf\(233.31 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Recent advances in wireless data networking and portable information appliances have engendered a new paradigm of computing, called mobile computing, in which users carrying portable devices have access to data and information services regardless of their physical location or movement behavior. In the meantime, research addressing information access in mobile environments has proliferated. In this survey, we provide a concrete framework and categorization of the various way ...

**Keywords:** application adaptation, cache invalidation, caching, client/server, data dissemination, disconnected operation, mobile applications, mobile client/server, mobile computing, mobile data, mobility awareness, survey, system application

### 2 [OceanStore: an architecture for global-scale persistent storage](#)



John Kubiawicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishna Gummadi, Sean Rhea, Hakim Weatherspoon, Chris Wells, Ben Zhao

November 2000 **ACM SIGARCH Computer Architecture News**, **ACM SIGOPS Operating Systems Review**, **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems ASPLOS-IX**, Volume 28, 34 Issue 5, 5


**Publisher:** ACM Press

Full text available: [pdf\(166.53 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers, data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks; monitoring also enhances performance through pro-active movement ...

### 3 [OceanStore: an architecture for global-scale persistent storage](#)

-  John Kubiawicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishan Gummadi, Sean Rhea, Hakim Weatherspoon, Westley Weimer, Chris Wells, Ben Zhao

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

**Publisher:** ACM Press

Full text available:  pdf(1.47 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers, data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks; monitoring also enhances performance through pro-active movement ...

4 Distributed systems - programming and management: On remote procedure call

Patrícia Gomes Soares


November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2**

**Publisher:** IBM Press

Full text available:  pdf(4.52 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)


The Remote Procedure Call (RPC) paradigm is reviewed. The concept is described, along with the backbone structure of the mechanisms that support it. An overview of works in supporting these mechanisms is discussed. Extensions to the paradigm that have been proposed to enlarge its suitability, are studied. The main contributions of this paper are a standard view and classification of RPC mechanisms according to different perspectives, and a snapshot of the paradigm in use today and of goals for t ...

5 Higher-order distributed objects

 Henry Cejtin, Suresh Jagannathan, Richard Kelsey

September 1995 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 17 Issue 5

**Publisher:** ACM Press

Full text available:  pdf(2.33 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We describe a distributed implementation of Scheme that permits efficient transmission of higher-order objects such as closures and continuations. The integration of distributed communication facilities within a higher-order programming language engenders a number of new abstractions and paradigms for distributed computing. Among these are user-specified load-balancing and migration policies for threads, incrementally linked distributed computations, and parameterized client-server applicat ...

**Keywords:** concurrency, continuations, higher-order languages, message-passing

6 File servers for network-based distributed systems

 Liba Svobodova

December 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 4

**Publisher:** ACM Press

Full text available:  pdf(4.23 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

7 A structural view of the Cedar programming environment

Daniel C. Swinehart, Polle T. Zellweger, Richard J. Beach, Robert B. Hagmann



August 1986 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,  
Volume 8 Issue 4

**Publisher:** ACM Press

Full text available: pdf(6.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents an overview of the Cedar programming environment, focusing on its overall structure—that is, the major components of Cedar and the way they are organized. Cedar supports the development of programs written in a single programming language, also called Cedar. Its primary purpose is to increase the productivity of programmers whose activities include experimental programming and the development of prototype software systems for a high-performance personal computer. T ...

## 8 Programming languages for distributed computing systems



Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum

September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

**Publisher:** ACM Press

Full text available: pdf(6.50 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

## 9 A survey of peer-to-peer content distribution technologies



Stephanos Androutsellis-Theotokis, Diomidis Spinellis

December 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 4

**Publisher:** ACM Press

Full text available: pdf(517.77 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance. Content distribution is an important peer-to-peer application ...

**Keywords:** Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

## 10 Special issue on prototypes of deductive database systems: The CORAL deductive system



Raghu Ramakrishnan, Divesh Srivastava, S. Sudarshan, Praveen Seshadri

April 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**,  
Volume 3 Issue 2

**Publisher:** Springer-Verlag New York, Inc.

Full text available: pdf(3.03 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

CORAL is a deductive system that supports a rich declarative language, and an interface to C++, which allows for a combination of declarative and imperative programming. A CORAL declarative program can be organized as a collection of interacting modules. CORAL supports a wide range of evaluation strategies, and automatically chooses an efficient strategy for each module in the program. Users can guide query optimization by selecting from a wide range of control choices. The CORAL system provides ...



**Keywords:** deductive database, logic programming system, query language

11 Rover: a toolkit for mobile information access



A. D. Joseph, A. F. de Lespinasse, J. A. Tauber, D. K. Gifford, M. F. Kaashoek  
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29 Issue 5

**Publisher:** ACM Press

Full text available: [pdf\(2.18 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 Design and development of data-intensive web sites: The Araneus approach



Paolo Merialdo, Paolo Atzeni, Giansalvatore Mecca  
February 2003 **ACM Transactions on Internet Technology (TOIT)**, Volume 3 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(2.18 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data-intensive Web sites are large sites based on a back-end database, with a fairly complex hypertext structure. The paper develops two main contributions: (a) a specific design methodology for data-intensive Web sites, composed of a set of steps and design transformations that lead from a conceptual specification of the domain of interest to the actual implementation of the site; (b) a tool called Homer, conceived to support the site design and implementation process, by allowing the ...

**Keywords:** Databases, Internet, WWW, World Wide Web, development

13 Interposed request routing for scalable network storage



February 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(363.12 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper explores interposed request routing in Slice, a new storage system architecture for high-speed networks incorporating network-attached block storage. Slice interposes a request switching filter---called a *μproxy*---along each client's network path to the storage service (e.g., in a network adapter or switch). The *μproxy* intercepts request traffic and distributes it across a server ensemble. We propose request routing schemes for I/O and file service traffic, and explore th ...

**Keywords:** Content switch, file server, network file system, network storage, request redirection, service virtualization

14 The family of concurrent logic programming languages



Ehud Shapiro  
September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

**Publisher:** ACM Press

Full text available: [pdf\(9.62 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Concurrent logic languages are high-level programming languages for parallel and distributed systems that offer a wide range of both known and novel concurrent programming techniques. Being logic programming languages, they preserve many

advantages of the abstract logic programming model, including the logical reading of programs and computations, the convenience of representing data structures with logical terms and manipulating them using unification, and the amenability to metaprogrammin ...

15 Middleware performance analysis: Performance monitoring of java applications



M. Harkema, D. Quartel, B. M. M. Gijsen, R. D. van der Mei

July 2002 **Proceedings of the 3rd international workshop on Software and performance WOSP '02**

**Publisher:** ACM Press

Full text available: pdf(219.69 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Over the past few years, Java has evolved into a mature platform for developing enterprise applications. A critical factor for the commercial success of these applications is end-to-end performance, e.g., in terms of response times, throughput and availability. This raises the need for the development, validation and analysis of performance models to predict performance metrics of interest. To develop and validate performance models, insight in the execution behavior of the application is essent ...

**Keywords:** performance measurement and monitoring of java applications

16 Introducing Ada 9X



John Barnes

November 1993 **ACM SIGAda Ada Letters**, Volume XIII Issue 6

**Publisher:** ACM Press

Full text available: pdf(4.39 MB) Additional Information: [full citation](#), [citations](#), [index terms](#)

17 Version models for software configuration management



Reidar Conradi, Bernhard Westfechtel

June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2

**Publisher:** ACM Press

Full text available: pdf(483.54 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

After more than 20 years of research and practice in software configuration management (SCM), constructing consistent configurations of versioned software products still remains a challenge. This article focuses on the version models underlying both commercial systems and research prototypes. It provides an overview and classification of different versioning paradigms and defines and relates fundamental concepts such as revisions, variants, configurations, and changes. In particular, we foc ...

**Keywords:** changes, configuration rules, configurations, revisions, variants, versions

18 Frontmatter (TOC, Letters, Election results, Software Reliability Resources!, Computing Curricula 2004 and the Software Engineering Volume SE2004, Software Reuse Research, ICSE 2005 Forward)



July 2005 **ACM SIGSOFT Software Engineering Notes**, Volume 30 Issue 4

**Publisher:** ACM Press

Full text available: pdf(6.19 MB) Additional Information: [full citation](#), [index terms](#)

## 19 Probabilistic object bases



Thomas Eiter, James J. Lu, Thomas Lukasiewicz, V. S. Subrahmanian

September 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 3

**Publisher:** ACM Press

Full text available: pdf(663.73 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Although there are many applications where an object-oriented data model is a good way of representing and querying data, current object database systems are unable to handle objects whose attributes are uncertain. In this article, we extend previous work by Kornatzky and Shimony to develop an algebra to handle object bases with uncertainty. We propose concepts of consistency for such object bases, together with an NP-completeness result, and classes of probabilistic object bases for which consi ...

**Keywords:** Consistency, object-oriented database, probabilistic object algebra, probabilistic object base, probability, query language, query optimization

## 20 The evolution of Coda



M. Satyanarayanan

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

**Publisher:** ACM Press

Full text available: pdf(441.35 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Failure-resilient, scalable, and secure read-write access to shared information by mobile and static users over wireless and wired networks is a fundamental computing challenge. In this article, we describe how the Coda file system has evolved to meet this challenge through the development of mechanisms for server replication, disconnected operation, adaptive use of weak connectivity, isolation-only transactions, translucent caching, and opportunistic exploitation of hardware surrogates. For eac ...

**Keywords:** Adaptation, Linux, UNIX, Windows, caching, conflict resolution, continuous data access, data staging, disaster recovery, disconnected operation, failure, high availability, hoarding, intermittent networks, isolation-only transactions, low-bandwidth networks, mobile computing, optimistic replica control, server replication, translucent cache management, weakly connected operation

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)